

AI GOVERNANCE FOR 2027

PREPARING
THE YEAR AHEAD

MODELS

AGENTS

DECISIONS



COGNAONE.COM

SEPTEMBER 2026



■ REGULATORY REVIEW · SEPTEMBER 2026 · V1.0

AI Governance: Preparing for 2027

Agents, automated decisions and the controls worth building now

Daniel Mackle · CognaOne · cognaone.com

A practical guide to the dated AI obligations of the next fifteen months across the EU, the United States, the UK, Singapore and Australia, the reporting clocks that already run, and the controls worth building now. It adds a tiering model for agents, a control matrix that serves several regimes at once, scenarios for the developments that are only signalled, and a plan from October 2026 to December 2027.

There is no single 2027 agent law to prepare for. There are several dated duties, each with its own scope, a set of reporting clocks that already run, and a growing expectation that enterprises can explain and constrain what their agents do.

CONTENTS

The useful question for 2027

1. What is actually coming

2. Obligations already running

3. The 2027 dividing line: consequential decisions

4. The EU deadline is a classification exercise first

5. Australia: a dated scam obligation and an active prudential expectation

6. What to monitor without inventing a deadline

7. Tier agents before you govern them

8. Build once, satisfy many

9. Plan for the signals with scenarios

10. Risk appetite and contracts

11. Who owns what

12. A control case for one agent

13. The plan: now through December 2027

14. What the board should ask for

Conclusion

About the author and CognaOne

References

The useful question for 2027

An agent does not become regulated because it is called an agent. The obligations turn on what the system does, whom it affects, where it operates, and whether the organisation develops it, deploys it or supplies a component. A research assistant and a system that decides whether to extend credit may use the same model and face very different requirements.

The next fifteen months bring several real dates. Before the year is out, Australia's automated-decision transparency rule commences and the EU's marking transition for existing generative systems ends. On 1 January 2027, California's rules for automated decisionmaking technology in significant decisions and Colorado's replacement AI law begin. Australia's Scams Prevention Framework moves into its main operating phase on 31 March. The EU's high-risk AI rules for Annex III systems apply from 2 December, and the main obligations of its Cyber Resilience Act follow on 11 December for products within that Act^{1,2,3,4,5}.

Other developments are signals rather than enacted 2027 deadlines: Australian AI legislation, Singapore's proposed financial-sector guidelines, UK agentic AI guidance and a statutory UK code, final EU classification guidance, and United States federal action against state AI laws^{6,7,8,9,10}.

An enterprise can prepare for the dated obligations without guessing the outcome of the proposals. The work begins with a register of systems and actions, a defensible applicability assessment, control at the point where an agent acts, and records that show what the system actually did. The legal requirements differ. The operational questions are remarkably similar.

The guide has two halves. Sections 1 to 6 set out what is dated, what is already running and what is only signalled. Sections 7 to 11 turn that into strategy: how to tier agents, which controls serve several regimes at once, how to plan for the signals, and what to put in risk appetite statements, contracts and the operating model. The tiering model, control matrix, scenarios and example wording are our analysis, offered as a starting point to adapt.

This is a forward-looking research and implementation guide, not a legal opinion or a claim that every measure below applies to every agent. Dates and scope are stated as at 28 September 2026. Recheck official texts before a release or market entry. It builds on our Global and Australian Focus reviews, which cover the instruments already in force in more depth.

1. What is actually coming

DATE OR PERIOD	JURISDICTION AND DEVELOPMENT	CURRENT STATUS	ENTERPRISE IMPACT
2 December 2026	EU AI Act: end of the transitional period for Article 50(2) marking of generative outputs by systems already on the market	Enacted	Providers of generative systems placed on the market before 2 August 2026 must have machine-readable marking of synthetic outputs in place ^{1,11} .

DATE OR PERIOD	JURISDICTION AND DEVELOPMENT	CURRENT STATUS	ENTERPRISE IMPACT
10 December 2026	Australia: APP 1.7 to 1.9 automated-decision transparency	Enacted, scheduled	In-scope APP entities must describe specified automated decisions and the personal information used in their privacy policies. The test also reaches a program doing something substantially and directly related to a significant decision ¹² .
December 2026	Australian Commonwealth: remaining requirements of the responsible-AI policy v2.0	Mandatory for covered non-corporate Commonwealth entities, phased	Government teams need to finish use-case registers, accountability and impact-assessment processes according to the policy's schedule. Suppliers should check contract terms rather than assume the policy binds them directly ¹³ .
1 January 2027	California: automated decisionmaking technology (ADMT) rules	Final regulations approved September 2025	Businesses using ADMT to make significant decisions must comply with the ADMT requirements, including pre-use notice, opt-out and access rights, from 1 January 2027. Risk-assessment obligations began earlier, with submissions to the California Privacy Protection Agency (CPPA) due by April 2028 ^{2,14} .
1 January 2027	Colorado: SB26-189 on automated decision-making technology	Signed May 2026; Attorney General rules due by 1 January 2027	Developers of technology used to materially influence consequential decisions must give deployers technical documentation. Deployers must give notice at the point of interaction, explain adverse outcomes within 30 days, and offer data correction and meaningful human review. Both keep records for three years ³ .
31 March 2027	Australia: main Scams Prevention Framework duties	Enacted framework with staged application	Regulated banking, telecommunications and digital-platform services must take steps to prevent, detect, disrupt, respond to and report scams. AFCA membership for covered entities began on 1 September 2026; further obligations are expected by the end of 2027 ⁴ .
2 December 2027	EU AI Act: high-risk Annex III systems	Enacted following the July 2026 amendment	Relevant providers and deployers need the Act's risk-management, documentation, logging, oversight, monitoring and other role-specific measures. The date does not make every AI agent high risk. Annex I product-system rules follow on 2 August 2028 ^{1,15} .
11 December 2027	EU Cyber Resilience Act (CRA): main obligations	Enacted	Manufacturers of in-scope products with digital elements face lifecycle cybersecurity and conformity duties. Reporting of actively exploited vulnerabilities and severe incidents already applies from 11 September 2026. A hosted agent service is not automatically a CRA product; classify the offering first ^{5,16} .

A date is not a compliance strategy. The first decision is applicability. For each agent, identify the legal entity, geography, affected people, use case, system role and route to market. Then determine which instruments attach. A single organisation may be a provider in one workflow, a deployer in another and a service provider under an institution's contract in both.

2. Obligations already running

Several duties that matter for agents are already in force, and 2027 preparation should not treat them as future events. In the EU, the AI Act's prohibitions, general-purpose model provisions and Article 50 transparency duties already apply, and the CRA's reporting obligations began on 11 September 2026^{15,16}. In the UK, the Data (Use and Access) Act 2025 reforms to automated decision-making have been in force since February 2026, and its remaining provisions since June 2026^{17,18}. In Australia, APRA's prudential standards, the Privacy Act's security and breach-notification duties, ransomware payment reporting and the Scams Prevention Framework's first stage all apply now; our Australian Focus edition covers them in detail^{19,20,21}.

The practical consequence is that an agent incident in 2026 can already start several reporting clocks, before any 2027 date arrives:

OBLIGATION	TRIGGER	CLOCK
EU CRA, Article 14	Actively exploited vulnerability or severe incident affecting an in-scope product	Early warning within 24 hours, notification within 72 hours, then a final report ¹⁶
APRA CPS 230	Operational risk incident with a material financial impact or material impact on a critical operation	Notify APRA within 72 hours ¹⁹
APRA CPS 234	Material information security incident	Notify APRA within 72 hours ²⁰
Australian Cyber Security Act 2024	Ransomware or cyber extortion payment by or on behalf of a reporting business entity	Report to ASD within 72 hours ²²
Australian Notifiable Data Breaches scheme	Suspected eligible data breach, likely to cause serious harm	Assessment within 30 days at most; notify as soon as practicable ²¹

Each clock needs the same thing: a record of what the agent accessed, proposed and did, fast enough to support a decision within hours. Most agent deployments do not yet produce that record by default.

3. The 2027 dividing line: consequential decisions

California and Colorado bring earlier 2027 obligations into focus for systems affecting people. They do not use identical definitions. California's regulations define ADMT as technology that processes personal information and uses computation to replace or substantially replace human decision-making; the CPPA says businesses using ADMT to make significant decisions must comply from 1 January 2027. Colorado's replacement law covers technology used to materially influence consequential decisions in areas including employment, housing, lending, insurance, healthcare and government services^{2,3}.

Neither regime is a checklist that a nominal human sign-off satisfies. Examine what information the reviewer receives, what they review and whether they can actually change the outcome. Colorado's Attorney General must adopt rules on adverse-outcome disclosures by 1 January 2027 and has proposed rules that include a definition of "materially influence"; those rules may arrive only weeks before the law applies^{3,23}. Colorado's law is also exposed to federal pressure: a December 2025 executive order directs federal challenges to state AI laws considered inconsistent with national policy¹⁰. Plan for the state obligation while expecting it to be contested.

Australia's APP 1.7 to 1.9 requirement starts earlier, on 10 December 2026. Its purpose is privacy-policy transparency, with a particular statutory trigger. It is not a general prohibition on automated decisions and does not, by itself, create California's or Colorado's individual rights¹².

The common preparation task is to find the workflows in which personal information is used to make, or substantially and directly support, decisions with serious consequences. A practical decision register should record:

- the decision and the person affected;
- the system's precise role: information retrieval, recommendation, ranking, approval or execution;
- the personal information and source systems used;
- where the individual, the enterprise and the system are located;
- who can review or reverse the outcome, and when;
- the notice, explanation, review and record-retention requirements found applicable by counsel.

Start with credit, employment, insurance, healthcare, access to essential services and public benefits. Do not classify every agent as a consequential-decision system simply because it helps a team work faster.

4. The EU deadline is a classification exercise first

The EU AI Act's high-risk obligations for Annex III systems begin on 2 December 2027. Annex III is use-case based. Employment, certain creditworthiness assessments, access to specified essential services, education, law enforcement and other listed areas require close review. There are classification qualifications and exceptions. The Commission published draft guidelines on high-risk classification in May 2026; final guidelines have not been dated⁹. An enterprise should document why a system is or is not high risk, including whether an agent is merely ancillary to a decision or performs a listed function. It should also distinguish provider obligations from deployer obligations and assess whether modifications could change its role¹⁵.

For an in-scope high-risk system, the preparation is substantial: define the intended purpose; obtain the provider's instructions and technical material; establish appropriate human oversight; test performance in the operating context; monitor it; retain the required automatically generated logs under the deployer's control; and agree how a serious issue is suspended and escalated. Article 26

is specific about deployer duties, including competent oversight and a minimum log-retention period, subject to other applicable law. A generic statement that "a human is in the loop" does not show what authority that human actually has²⁴.

The CRA is a separate product-security question. It matters to an organisation putting in-scope software or hardware products with digital elements on the EU market, including remote data processing that is essential to a product's function. It requires a product classification and a manufacturer and supply-chain analysis. For a covered agent product, vulnerability handling, secure development, documentation and conformity work should already be underway, and the Article 14 reporting obligations already apply. Products placed on the market before 11 December 2027 come under the main obligations only if substantially modified after that date^{16,25}. The CRA should not be added mechanically to a cloud-only agent inventory without checking scope.

5. Australia: a dated scam obligation and an active prudential expectation

For covered Australian banks, telecommunications services and digital platforms, 31 March 2027 is an operational date. A payment or customer-communication agent can assist scam detection, but it can also be manipulated into authorising a payment, changing a contact destination or giving misleading reassurance. The Scams Prevention Framework applies to the designated service, not to the agent. Teams should map each agent action that can affect prevention, detection, disruption, response, reporting or a customer's complaint path, and check the sector rules as they are finalised⁴.

The Australian prudential position does not wait for 2027. APRA's April 2026 AI letter identifies weak agent identity and access management, reliance on policy rather than preventive technical controls, opaque supplier dependencies, inadequate fallback for AI used in critical operations and assurance that lags deployment. CPS 230 and CPS 234 remain the binding operational-risk and information-security foundations. APRA says it is finalising a forward plan for AI supervision, including prudential reviews and supplier engagement, but has given no date for a separate AI standard. Treat the letter as a supervisory signal about existing duties, not new legislation^{19,20,26}.

The Government said in July 2026 that it expected to legislate Australian Standards for AI in early 2027, focused on large data centres and AI training. That is a policy intention, not an enacted agent-deployment regime. It warrants monitoring for infrastructure and model-training activities, but it is a poor basis for telling every Australian agent team that a general AI Act is arriving in 2027⁶.

6. What to monitor without inventing a deadline

Singapore's MAS consultation P017-2025 proposes AI risk-management guidelines for financial institutions. The consultation closed in January 2026 and the guidelines remain proposed. IMDA's model framework for agentic AI offers voluntary design guidance on bounding autonomy and action space, human accountability, technical controls and testing^{7,27}.

In the UK, the ICO has said its 2026/27 work includes an AI code of practice and dedicated guidance on agentic AI. The FCA's Mills Review addresses the move toward more autonomous retail finance and builds on existing outcomes-based obligations, including the Consumer Duty and the Senior Managers Regime. Neither establishes a general new UK agent law with a fixed 2027 date^{8,28}.

DEVELOPMENT	WHAT WOULD CHANGE THE ASSESSMENT	WORK WORTH DOING NOW
EU classification guidelines and standards	Final Commission guidelines on high-risk classification; harmonised standards	Classify Annex III candidates against the draft guidelines and record the reasoning; revisit when the final text appears ⁹
Colorado Attorney General rules	Final rules on adverse-outcome disclosures and "materially influence"	Inventory tools that score, rank or recommend in consequential decisions; draft notices that can be adjusted to the final rules ^{3,23}
US federal action on state AI laws	Litigation or federal measures affecting Colorado or California	Keep state compliance on track; track outcomes rather than assuming pre-emption ¹⁰
MAS P017-2025	Final guidelines, commencement and any transition	Map AI use, materiality, third-party dependencies and incident ownership against the consultation ⁷
UK ICO and FCA	Final ICO code or agentic guidance; FCA rule changes or supervisory findings	Test significant-decision review, data provenance, customer outcomes and complaint escalation against existing law ^{8,28}
Australian Standards for AI	A bill, enacted scope and commencement	For data-centre or training activities, assess exposure; for ordinary deployment, continue with existing privacy, prudential, cyber and sector duties ⁶
Australian SPF sector detail	Final rules, codes and further obligations	Assign a service owner, rehearse scam intervention and complaints, and update controls when the instruments settle ⁴

These are monitoring items with an owner and a review date, not placeholders in a compliance calendar labelled "deadline".

7. Tier agents before you govern them

Treating every agent the same either over-governs the harmless ones or under-governs the dangerous ones. Governance intensity should follow what an agent can do, not what it is called. The four tiers below are our model. They sit alongside legal classification rather than replacing it: a Tier 2 agent can still fall in an EU Annex III category, and a Tier 4 agent may be outside every regime in Section 1 and still need strong controls because of what it can move.

TIER	WHAT THE AGENT CAN DO	MINIMUM CONTROLS	APPROVAL TO DEPLOY	REVIEW AND EVIDENCE
1. Assist	Reads, retrieves and summarises for a person. Takes no action outside the conversation.	Register entry with owner; approved data sources; no access to sensitive data unless justified; staff guidance on use	Business owner	Annual review; usage and data-access logs
2. Recommend	Drafts, scores, ranks or recommends. A person decides and acts.	Tier 1, plus documented purpose and limits; testing for quality and bias where people are affected; a reviewer who sees the inputs and can reject the output	Business owner and risk	Six-monthly review; sample checks of recommendations against outcomes
3. Act within limits	Writes records, sends messages or triggers workflows inside defined limits, without per-action approval.	Tier 2, plus a distinct agent identity with least privilege; hard limits on value, volume, recipients and destinations; pre-execution policy and state checks; hold and escalate when uncertain; attributable action logs; tested fallback	Risk and security sign-off	Quarterly review; denied and held actions reported; incident rehearsal
4. Act on consequential matters	Makes or materially influences decisions about people, moves value, or acts in a critical operation.	Tier 3, plus legal classification against each applicable regime; human approval at defined high-impact checkpoints by a reviewer with real authority; notices, explanation and review routes where required; decision records kept for the longest applicable period; supplier exit plan	Executive owner and risk committee	Continuous monitoring; independent assurance of the operating workflow; board reporting

Two rules keep the model honest. An agent's tier is set by its most consequential permission, not its usual behaviour: an agent that normally summarises but can send payments is Tier 4. And a tier change is a change event, reviewed before new permissions are granted, which is where APRA's concern about change and release control meets agents directly²⁶.

8. Build once, satisfy many

The regimes in this guide differ in scope and wording, but they keep asking for the same few capabilities. Building those once, well, is cheaper and more defensible than running a separate programme per jurisdiction. The table shows which obligations or expectations each control helps meet. It indicates where a control provides evidence, not that the control alone achieves compliance.

CONTROL	WHAT IT HELPS MEET
Agent and use-case register with owner, purpose, tier, data and permissions	APRA's expected inventory of AI tooling and use cases ²⁶ ; the Commonwealth use-case register ¹³ ; knowing which programs make or support decisions for APP 1.7 ¹² ; identifying covered technology for California and Colorado ^{2,3} ; recording EU classification reasoning ¹⁵
Distinct identity and least privilege for each agent	CPS 234 information security capability ²⁰ ; APRA's finding that access management has not adjusted to AI agents ²⁶ ; APP 11 security of personal information the agent can reach ²⁹
Pre-execution checks and holds on consequential actions	APRA's expectation of preventative rather than detective controls ²⁶ ; staying within the risk appetite a Board has approved under CPS 220 or SPS 220 ^{30,31}
Attributable decision and action records	EU deployer log retention for high-risk systems ²⁴ ; Colorado's three-year records ³ ; every reporting clock in Section 2 ^{16,19,20,21,22}
Meaningful human review by someone with authority to change the outcome	EU human oversight for high-risk systems ²⁴ ; Colorado's right to request meaningful human review ³ ; whether a system replaces human decision-making is part of California's ADMT test ²
Notices and explanations	California pre-use notice and access rights ^{2,14} ; Colorado pre-use and adverse-outcome disclosures ³ ; APP 1.7 to 1.9 privacy-policy content ¹²
Fallback and supplier management	CPS 230 critical operations and service providers ¹⁹ ; APRA's findings on supplier concentration and contracts ²⁶ ; CRA supply-chain duties for manufacturers of covered products ²⁵

The strategic choice this frames is between one global standard with local overlays and separate programmes per market. For most enterprises the first is better: the seven controls above form the global baseline, and jurisdiction-specific notices, rights handling and retention periods sit on top. The exception is a business whose regulated activity is concentrated in one market, where a local programme may be simpler.

9. Plan for the signals with scenarios

A watchlist tells you what to monitor. A scenario tells you what to do. For each signalled development below, the no-regret moves are worth doing whichever way it goes; the held moves wait for the trigger.

SCENARIO	TRIGGER TO WATCH	NO-REGRET MOVES NOW	MOVES TO HOLD UNTIL TRIGGERED
EU dates or standards slip again	Further amendment of the AI Act timetable; harmonised standards still not published through 2027. The high-risk dates have already moved once ¹	Classify systems and keep the register current; build logging and oversight that are useful anyway	Formal conformity documentation and supplier certification demands tied to specific standards

SCENARIO	TRIGGER TO WATCH	NO-REGRET MOVES NOW	MOVES TO HOLD UNTIL TRIGGERED
US federal challenge to state AI laws succeeds or fails	Court rulings or federal measures under the December 2025 executive order ¹⁰	Meet California and Colorado duties on their current dates; keep notices and review processes modular	Retiring state-specific processes, or extending them to other states
Australia widens its approach	A bill for Australian Standards for AI whose scope reaches deployment rather than data centres and training ⁶	Operate under existing prudential, privacy, cyber and scams duties; keep the register and tiering in place	Any new registration, assessment or certification process the bill might require
Financial-sector guidance is finalised	Final MAS guidelines; APRA's forward plan for AI supervision; ICO agentic guidance ^{7,8,26}	Run the consultation drafts as a gap assessment; fix identity, supplier and assurance gaps APRA has already named	Re-mapping to the final text and any sector-specific templates

The pattern across all four is the same. The controls in Section 8 are no-regret in every scenario. What varies is paperwork, certification and jurisdiction-specific process, which can wait for the text.

10. Risk appetite and contracts

Two short instruments turn governance intent into something enforceable: the board's risk appetite and the supplier contract.

Example appetite statements for agents. These are illustrations to adapt, not recommended limits:

- No agent may initiate a payment, or change payment details, to a destination that has not been independently verified.
- No agent may make an adverse decision about an individual without review by a person with authority to change it.
- No agent may modify its own permissions or delegate authority to another agent without a recorded grant.
- An agent supporting a critical operation must have a fallback that keeps the operation within tolerance if the agent or its provider fails.
- Agent-caused incidents must be detectable and reconstructable within the shortest applicable reporting clock.

Supplier contract checklist. APRA found contracts with AI providers often lagged practice, with limited provisions on audit rights, model updates and deviations, incident notification or changes to data handling²⁶. A contract for a model, agent platform or embedded agent should address:

- notice before material model or behaviour changes, and the right to test before they apply;
- incident and vulnerability notification within a period that lets you meet your own reporting clocks;

- audit and information rights, including access to logs of the agent's actions on your behalf;
- where data is processed, whether it is retained or used for training, and how it is deleted;
- the documentation you need as a deployer, such as Colorado's developer documentation or an EU provider's instructions for use^{3,24};
- subcontracting, concentration and exit, including data return and a transition period.

11. Who owns what

Agent governance fails most often at the handoffs. A simple allocation that works in most enterprises:

FUNCTION	OWNS
Business owner	The agent's purpose, tier, outcomes and register entry; first-line controls
Technology and security	Agent identity, permissions, limits, pre-execution checks, logging and fallback
Legal and privacy	Applicability decisions, notices, rights handling and retention periods
Procurement	Supplier due diligence and the contract terms in Section 10
Risk (second line)	The tiering model, risk appetite, challenge and aggregate reporting
Internal audit (third line)	Independent testing of the operating workflow, not only the policy
Board and risk committee	Appetite, Tier 4 approvals and oversight of the programme

Voluntary standards can give this structure a recognised shape. ISO/IEC 42001 specifies an AI management system, and the NIST AI Risk Management Framework offers a widely used vocabulary for mapping, measuring and managing AI risk^{32,33}. Neither creates a legal obligation. Both are useful evidence that governance is systematic, particularly with customers and supervisors outside a single jurisdiction.

12. A control case for one agent

Consider a credit application agent used by a multinational lender. It retrieves an applicant's information, proposes a decision, drafts the notice and can update the customer record. The same workflow can raise different questions in each market: whether its EU use falls in an Annex III category; whether California's significant-decision ADMT provisions apply; whether Colorado's covered-decision provisions apply; and whether an Australian APP entity has the arrangements described in APP 1.7. This example illustrates a scoping method. It does not assume a single deployment is subject to all four regimes^{2,3,12,15}.

Before production, the lender should be able to answer:

1. **What is the agent allowed to change?** Reading a file, proposing a score, declining credit and sending an adverse notice are separate permissions.
2. **What evidence supports the decision?** Identify the input sources, their freshness, contested records and any correction route. Colorado and California both give individuals routes to correct or access the data used.
3. **Where does a person have authority?** If review is required, the reviewer should see the material facts and proposed action, and be able to change or stop it before the relevant consequence.
4. **What is captured?** Preserve the model and policy versions, state used, proposed and executed actions, approval or override, customer notice and actual outcome, under a retention schedule that meets the longest applicable requirement, such as Colorado's three years.
5. **What happens on failure?** An unavailable provider, policy-service outage, missing evidence or inconsistent applicant state needs a defined hold or fallback path.

An action gate is one possible engineering response: it checks authority, scope, policy and relevant state before a write or external action. The laws above do not prescribe one named product or require every action to be human-approved. The enterprise must choose controls that satisfy the obligations applicable to the particular workflow and can demonstrate that they work.

13. The plan: now through December 2027

The plan below is designed for an enterprise that already has agents in pilots or production. A smaller organisation can scale the formality, but should keep the decision points and evidence.

PERIOD	WORK TO COMPLETE	ACCOUNTABLE OWNER AND EVIDENCE OF COMPLETION
October 2026: find and classify	Inventory deployed and shadow agents, including agents embedded in purchased software. Capture owner, purpose, models, tools, credentials, data, countries, users, suppliers and every external or write action. Assign each agent a tier (Section 7), separating advisory systems from systems affecting consequential decisions or critical operations. Screen EU Annex III, California, Colorado, Australian APP and SPF scope. Confirm which reporting clocks already apply.	Business owner and legal or privacy lead sign an applicability record for each material workflow. Inventory reconciled against identity, API and procurement records.
November 2026: bound execution	Adopt the tiering model and risk appetite statements. Remove shared human credentials. Give each agent a distinct identity and minimum necessary tool and data access. Define transaction, recipient, destination and rate limits. Place high-impact writes behind a policy check or approval. Specify hold and escalation behaviour for missing evidence, conflicting state and unavailable controls.	Security and platform owners provide effective-permission exports, policy versions and passing tests for blocked actions.

PERIOD	WORK TO COMPLETE	ACCOUNTABLE OWNER AND EVIDENCE OF COMPLETION
December 2026: finish immediate notices and evidence	EU generative-system providers confirm Article 50(2) marking before 2 December. Australian APP entities review privacy-policy content against APP 1.7 to 1.9 before 10 December. Commonwealth agencies finish their policy milestones. California and Colorado teams validate classifications, notices, documentation, rights handling and human-review procedures, and update Colorado material once the Attorney General's rules are final. Rehearse one end-to-end incident against the reporting clocks and one provider outage.	Privacy and legal approve notices; operations retain a complete decision trace, incident exercise and fallback test.
January to March 2027: operate the new US and Australian duties	Start renegotiating material AI supplier contracts against the checklist in Section 10. Run California and Colorado processes in production where applicable. Sample actual adverse decisions, notices, correction and review requests. For Australian designated services, finish scam detection, disruption, reporting, complaints and third-party handoffs ahead of 31 March.	Compliance and customer-operations owners report failed controls, review turnaround, customer outcomes and remediation to the risk committee.
April to September 2027: validate EU readiness	For classified EU high-risk systems, complete provider and deployer documentation, data and performance evaluation, human-oversight design, logging, monitoring and incident arrangements. For CRA-covered products, close secure-development and conformity gaps. Test version changes, prompt injection, stale data, cross-agent delegation and provider concentration.	EU product and legal owners maintain an evidence file for each in-scope system. Independent risk or audit tests the operating workflow rather than the policy document alone.
October to December 2027: release gates	Resolve open findings before the 2 and 11 December EU dates. Recheck legal classification, final guidance and third-party terms. Prevent an in-scope release if required evidence or controls are absent; record the decision and owner.	Release committee receives a system-specific readiness decision, test results, residual risks and an incident contact tree.

The first two months matter most. If an organisation cannot establish which agents exist or what each can execute, it cannot reliably classify the legal obligations or test the controls. A spreadsheet inventory may be enough at first. What matters is that it is complete, owned and reconciled to what actually runs.

14. What the board should ask for

A board does not need a demonstration of every model. It needs evidence that the organisation can govern the actions with consequences. A useful quarterly pack would show:

- number of agents in each tier, with owner and jurisdiction, and any tier changes;
- number that can write records, communicate externally, move value or materially influence decisions about people;
- unowned agents, shared credentials and exceptions to access policy;
- high-risk or consequential-decision classifications and the reason for each conclusion;
- denied actions, human escalations, reversals and incidents, with trends and root causes;
- whether the organisation could have met each applicable reporting clock for the incidents it had;

- supplier concentration and tested fallback for critical operations;
- upcoming obligations, accountable owners and unresolved readiness gaps;
- status of the scenarios in Section 9, and whether any trigger has been met.

These are management measures, not a statutory reporting template. They let a board challenge whether policy, permissions and actual behaviour agree. APRA's 2026 observations show why that distinction matters in regulated financial services²⁶.

Conclusion

There is no single 2027 agent law to prepare for. There are several dated duties, each with its own scope, a set of reporting clocks that already run, and a growing supervisory expectation that enterprises can explain and constrain what their agents do. The temptation is to begin with a universal control framework. The better approach is to build the few controls every regime asks for, tier agents so effort follows risk, and then take one material workflow at a time: identify its role and jurisdictions, follow an action from instruction to outcome, and test whether the organisation can stop it, explain it and recover from it.

Do that across the agent estate during the rest of 2026. By early 2027, the organisation should have a working register, decisions on applicability, production controls for consequential actions, and evidence from live operation. Those assets remain useful even where proposed guidance changes or legislation arrives later than expected.

About the author and CognaOne

Daniel Mackle is the founder of CognaOne. CognaOne is building the active governance platform for AI in regulated financial services: an inventory of AI systems and the agents connected to them, controls held with their regulatory provenance, an action gate that allows, holds or declines proposed actions and returns the same decision on the same facts and policy, and decision records created at the moment of authorisation. This review was written to stand on its sources, not on the product, and every factual statement in it is referenced to the issuing body wherever that body's text was accessible.

AI use: research and drafting for this review were supported by AI tools. Multiple models were used for research and to cross-check facts, and every factual statement was verified by the author against the cited sources. The author is responsible for the content.

References